

Visualization of Coregistered Imagery for Remote Surface Operations

Mark W. Powell, Paul G. Backes, Marsette A. Vona, Jeffrey S. Norris
Jet Propulsion Laboratory, California Institute of Technology, Pasadena, California
Mark.Powell@jpl.nasa.gov

Abstract—This work discusses various techniques used in visualizing sets of coregistered images for planetary lander or rover mission operations. Imagery from all available instrumentation on board a planetary rover or lander is useful in science analysis of a remote landing site. We use a visualization technique called data fusion to expedite analyzing data from an instrument suite. Data fusion uses coregistration information to overlay images from different instruments together. These overlay images are then presented to the scientist in both 2D and 3D visualization tools. Next, we introduce a technique for panorama visualization that is based on a spherical warp operation. This warping technique is optimized for computational efficiency in order to run dynamically and to support a 2D or 3D visualization tool. We then present some results of the implementation of these techniques for the 2003 MER Science Activity Planner software, using image data from the 2002 FIDO-MER field test. We conclude with a discussion of future directions for these visualization techniques.

TABLE OF CONTENTS

- 1 INTRODUCTION
- 2 COREGISTRATION VISUALIZATION
- 3 IMAGE MOSAICS
- 4 CONCLUSIONS
- 5 ACKNOWLEDGEMENTS

1. INTRODUCTION

The process of tactical science planning in the day-to-day operations of a major remote science mission is extremely demanding. New images and rover localization information become available at different times throughout each day, due to the need for multiple data transmission paths as the lines-of-sight to each of the cooperating communication mechanisms appear and disappear. For example, in the upcoming 2003 MER mission, the rovers each have a direct to Earth (DTE) antenna for transmitting information while they are on a line of sight to Earth, and a UHF antenna for transmitting data via the Mars Odyssey orbiter when the rover is occluded from Earth by Mars. The limited available bandwidth and complicated communication paths create the need for a mission operations schedule that is largely reactive. The time available for making tactical science planning decisions is therefore very limited, and quick assessment of rover localization

and visualization of recently available data products is essential. Software tools that provide a science planning interface for mission operators must be very fast to be effective under these conditions.

During the 2002 FIDO-MER field test, the pace was even more frantic than that of an actual mission since the remote operations were run at a compressed schedule where 2 sols of operations were planned and executed each day. This compressed schedule provided only 45 minutes at the most (and often less due to communication delays) to assess the opportunities to conduct science based on the newly available images and other telemetry before the participants had to devote their complete attention to generating a plan for the rover to execute. For operations under these kinds of serious time constraints, high-performance visualization tools are essential to maximizing the science return of a mission.

The following sections describe two advanced visualization techniques for visualizing the coregistered images that are captured in situ by a remote rover quickly and effectively. The first is data fusion, which allows the visualization of images from different instruments or images taken from different points of view to be rendered as a overlaid projection. This allows the analyst to immediately see the qualitative coregistration of multiple images to help to assess the quality of these data products. The next visualization technique is a spherical panorama warp operation, which we show is ideally suited for displaying a panorama of coregistered images that were captured by a camera on a pan/tilt mechanism such as a rover camera mast. We conclude with a brief summary and a discussion of future directions for this line of visualization research.

2. COREGISTRATION VISUALIZATION

Visualization of coregistered imagery (first presented in [1]) is useful in conducting remote science. When multiple cameras with different reflectance filters take images of a given area, it is useful to see an image of that area that is comprised of reflectance data from all of the instruments combined. When a given area is imaged from different points of view by a rover using the same instrument, it is useful to visualize the coregistration of these images as a measure of the accuracy of rover localization.

During the 2002 FIDO-MER field test [2], the accuracy of the 3D information computed from the front hazard camera (Hazcam) stereo images was called into question based on some

anomalous results from instrument arm activities that relied on that data. In order to diagnose the problem, scientist participants manually correlated pixels of areas imaged by both the mast-mounted navigation camera (Navcam) and the Hazcam. This laborious process enabled them to better quantify the errors in the Hazcam range data. This took several days to accomplish, and furthermore the results were only valid for the rover configuration at that particular time. If the cameras or other mechanisms had exhibited the tendency to change over time (for instance, as the result of thermal variations in the optical bench, or an impact or other physical degradation that altered the effective kinematics of the mechanism), the results would need to be recomputed. Visualizing this coregistration automatically by projecting the 3D data from the Hazcam into the Navcam image makes this inter-instrument correlation procedure simple to accomplish as a mission progresses, perhaps during every sol of operations if necessary.

Data Fusion

Our method of projecting the information from one image onto another image is called data fusion. This method will work for any two images where one image has a registered xyz map (such as those produced by means of stereo image correlation or laser range finder data) and the other image has an associated camera model that can relate xyz points in the world to 2D image coordinates. For our implementation of the method, we use camera models defined as CAHV [3], CAHVOR[4], or CAHVORE. These camera model specifications are particularly well-suited for performing 3D to 2D and 2D to 3D projections easily. The general idea of the process is illustrated in Figure 1. Image A has a registered XYZ map, and the pixel at the tail of the yellow arrow emanating from A corresponds to the 3D point at the arrow head. This point projects through a camera model associated with Image B onto a pixel in B. For any pixel in image A the xyz value for that pixel is projected to a particular set of 2D image coordinates in image B. If the coordinates lay within the boundary of image B, the color of the pixel at those coordinates is then overlaid onto image A. We can produce a complete image overlay by iterating over every pixel in image A, building up an image of projected pixels from image B copied to the corresponding pixel positions in image A.

Figure 2(c) shows an example of one image projected and overlaid onto another. The background image is a FIDO Navcam image and the center, bright region is an overlaid Pancam image. Although the images appear to be approximately overlaid in the correct positions, they are not overlaid perfectly. This is particularly evident by comparing the actual position of the bright feature that lies just above the upper right portion of the image overlay to the projection of that feature from the Pancam image. Since these images were taken from the same rover position, these errors could be due to a number of other factors, such as physical positioning of the camera pan/tilt mast, or the fidelity of the camera modeling and calibration.





Figure 1. Example of projection of a pixel from one image to another. Image A has a registered XYZ map, and the pixel at the tail of the yellow arrow emanating from A corresponds to the 3D point at the arrow head. The second arrow shows the projection of this point onto a pixel in Image B.

3. IMAGE MOSAICS

Viewing a collection of images as a mosaic or a panorama is a valuable tool for scientists conducting remote rover operations. Viewing panoramas taken by a rover between major sites provides an effective localization mechanism. A mosaic is useful both for strategic planning (such as determining the safest general route to follow for a long drive) and tactical planning (prioritizing the best target areas for remote observations such as spectroscopy). Creating panoramas involves complicated camera modeling and image projection techniques, making it an expensive and time-consuming procedure. The daily mission operation schedule for remote science demands planning tools that are maximally time-efficient, and so we have devised methods of computing mosaic images dynamically that are highly optimized for speed while still producing high-quality images.


For the Mars Pathfinder mission, the very popular cylindrically projected mosaic was produced from images taken from the lander. This panorama was projected cylindrically since the lander's camera acquired images through 360 degrees of azimuth but a limited range of elevation. For the FIDO rover and the MER mission, the mast cameras are mounted on pan-tilt mechanisms that are capable of rotating through a full



(a)



(b)



(c)

Figure 2. Example of a FIDO Pancam image projected and overlaid onto a Navcam image. a) The Navcam background image. b) The Pancam foreground image. c) The Pancam projected and overlaid onto the Navcam.

range of azimuth and elevation. Furthermore, in the Athena science team for the MER mission, there has been a recent surge of interest in acquiring in-situ atmospheric imagery, for which a cylindrical projection is not an adequate visualization. We therefore chose a spherical projection as our preferred method of mosaic construction.

Figure 3 shows an example of a spherically projected panorama using data from the 2002 FIDO-MER field test.

Warping algorithm

We will now describe in detail our method for spherical projection of a set of images taken by a camera on a pan-tilt mechanism. The sphere that the images are projected onto is centered at a point that is at the center of rotation of the pan-tilt mechanism. The radius of the sphere can be set almost arbitrarily, but a good rule of thumb is that the radius should be the height of the sphere center from the ground at a minimum. In practice, larger radii that differ from this distance by a factor of 10 or 100 had no noticeable effect on the resulting warped image.

For each image, there must be an associated camera model. The minimum requirements of this model are that the 3D point that is the center of projection of the optical path of the camera is known, and is (or can be made) relative to the same coordinate frame as the position of the sphere center. Additionally, the camera model must relate each pixel position on the 2D image plane (i, j) with a 3D ray into the world into which this pixel projects. Given this ray, the intersection with the sphere is computed, and the resulting 3D intersection point (x, y, z) is converted to spherical coordinates (θ, ϕ, ρ) . If this intersection is computed for every pixel in every image, we can create a 2D raster where the horizontal axis is θ and the vertical axis is ϕ . An example of the resulting image is shown in Figure 3.

The procedure above for creating the mosaic is sufficient but can be made more efficient. Since the camera model for each image must be able to derive the 3D ray that projected onto any 2D image plane coordinates, we can also reverse the association and derive the 2D image position from any ray that intersects the image plane. Back projecting rays from the sphere onto the source images can result in a substantial time savings for several reasons. First, depending on the desired resolution of the warped image, adjacent pixels in a particular source image may project onto the same area of the sphere, so by projecting from the sphere onto the images, the number of intersections to compute is reduced from many to one. Also, since the individual images are purposefully taken such that adjacent images overlap in their field of view, these large areas of overlap result in redundant computations if the projection from the image to the sphere is performed for every pixel in every image.

Given that the method of projection from the sphere onto an image to create the mosaic is preferred, there is the problem




Figure 3. Spherically projected mosaic of site 1 of the 2002 FIDO-MER field test.

of deciding which image of the set of images to project onto in order to form each pixel. The brute force solution is to simply attempt to intersect each ray from the sphere with each of the images iteratively, until the 2D image plane intersection lies within the boundary of the image. In order to optimize this search, one could compute the view frustum of the field of view of each image and iteratively test whether the ray is contained within the subtended volume. An even faster solution is this: First, precompute the (θ_i, ϕ_i) coordinates of the ray that is projected from the center pixel each source image. Iterate over each (θ_s, ϕ_s) of the sphere to project the ray onto an image, each time selecting the image having the minimum distance from the point on the sphere to the point intersected by the image direction of projection on the sphere, or:

$$\min(\sqrt{(\theta_s - \theta_i)^2 + (\phi_s - \phi_i)^2}) \quad (1)$$

Visualization of Mosaics

Since a spherically-projected mosaic can sometimes warp the images in such a way that is not aesthetically pleasing, we have devised a visualization tool that is ideally suited to viewing these images. The visualization is based on a 3D rendering system (for our implementation we chose Java3D). The viewpoint is placed at the center of a sphere and allows the user to rotate the viewpoint to any angle. The spherically projected mosaic is then texture mapped onto the interior of the sphere, creating a perspective projection image of whatever area of the sphere is currently visible (see Figure 4). This method of visualization is as pleasing to use as it is fast, since it takes advantage of hardware-accelerated texture mapping support on commonly available graphics hardware.

4. CONCLUSIONS

Coregistration visualization tools provide a significant savings in time and an effective means of qualitatively assessing

coregistration. Future directions of this research will address methods for quantitative analysis of coregistration errors.

The spherically warped panorama is ideally suited for visualized imagery taken from cameras on a pan/tilt mechanism. For future missions, we expect that orbital data sets and in-situ imagery will need to be coregistered and visualized, so future work in this area of research will address combining projections of imagery from orbiters with projections of images from in-situ rovers.


The coregistration visualization techniques described here are being implemented as a part of the Web Interface for Tele-science (WITS) remote operations software for remote science field tests and missions. The goals of the WITS project are to provide a superior collaborative science planning tool for mission participants in both the scientific and engineering areas while also making the mission process accessible to the public by creating a publicly-available version of the software for use on home computers or in schools. The WITS software that was used during the 2002 FIDO-MER field test is available on the WITS web site [5].

5. ACKNOWLEDGEMENTS

The research described in this paper was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

REFERENCES

- [1] M. Powell, J. Norris, and P. Backes, "Visualization of spectroscopy for remote surface operations," in *IEEE Aerospace Conference*, (Big Sky, MT), March 2002.
- [2] "<http://fido.jpl.nasa.gov>."
- [3] Y. Yakimovsky and R. Cunningham, "A system for extracting three-dimensional measurements from a stereo



(a)



(b)

Figure 4. Spherical projection of site 1 of the 2002 FIDO-MER field test as viewed in the 3D-accelerated spherical visualization tool.

pair of tv cameras,” *Computer Graphics and Image Processing*, vol. 7, pp. 195–210, 1978.

- [4] D. Gennery, “Camera calibration including lens distortion,” Tech. Rep. JPL D-8580, Jet Propulsion Laboratory, California Institute of Technology, 1991.


- [5] “<http://wits.jpl.nasa.gov>.”



Marsette A. Vona, III Marsette Vona is a computer scientist, software engineer, and electromechanical hardware developer. He is a member of the technical staff of the Mobility Systems Concept Development Section at the Jet Propulsion Laboratory. At JPL, his work is currently focused in the areas of high-performance interactive 3D data visualization for planetary exploration, user-interface design for science data analysis software, and Java software architecture for large, resource-intensive applications. Marsette received a B.A. in 1999 from Dartmouth College in Computer Science and Engineering, where he developed new types of hardware and algorithms for self-reconfigurable modular robots. He completed his M.S. in 2001 at MIT’s Precision Motion Control lab, where he developed a high-resolution interferometric metrology system for a new type of robotic grinding machine. Marsette was awarded the Computing Research Association Outstanding Undergraduate Award in 1999 for his research in Self-Reconfigurable Robotics.



Jeffrey S. Norris Jeff Norris is a computer scientist and member of the technical staff of the Mobility Systems Concept Development Section at the Jet Propulsion Laboratory. At JPL, his work is focused in the areas of distributed operations for Mars rovers and landers, secure data distribution, and science data visualization. Currently, he is a software engineer on the Mars Exploration Rover ground data systems and mission operation systems teams. Jeff received his Bachelor’s and Master’s degrees in Electrical Engineering and Computer Science from MIT. While an undergraduate, he worked at the MIT Media Laboratory on data visualization and media transport protocols. He completed his Master’s thesis on face detection and recognition at the MIT Artificial Intelligence Laboratory. He lives with his wife, Kamala, in La Crescenta, California.



Mark W. Powell Mark Powell is a member of the technical staff in the Mobility Systems Concept Development Section at the Jet Propulsion Laboratory, Pasadena, CA since 2001. He received his B.S.C.S. in 1992, M.S.C.S in 1997, and Ph.D. in Computer Science and Engineering in 2000 from the University of South Florida, Tampa. His dissertation work was in the area of advanced illumination modeling, color and range image processing applied to robotics and medical imaging. At JPL his area of focus is science data visualization and science planning for telerobotics. He is currently serving as a software and systems engineer, contributing to the development of science planning software for the 2003 Mars Exploration Rovers and the JPL Mars Technology Program Field Integrated Design and Operations (FIDO) rover task. He, his wife Nina, and daughters Gwendolyn and Jacquelyn live in Tujunga, CA.



Paul G. Backes Paul Backes is a technical group leader in the Mobility Systems Concept Development section at the Jet Propulsion Laboratory, Pasadena, CA, where he has been since 1987. He received the BSME degree from U.C. Berkeley in 1982, and MSME in 1984 and Ph.D. in 1987 in Mechanical Engineering from Purdue University. He is currently responsible for distributed operations research for Mars lander and rover missions at JPL. Dr. Backes received the 1993 NASA Exceptional Engineering Achievement Medal for his contributions to space telerobotics (one of thirteen throughout NASA), 1993 Space Station Award of Merit, Best Paper Award at the 1994 World Automation Congress, 1995 JPL Technology and Applications Program Exceptional Service Award, 1998 JPL Award for Excellence and 1998 Sole Runner-up NASA Software of the Year Award. He has served as an Associate Editor of the IEEE Robotics and Automation Society Magazine. Paul initiated the development of WITS/SAP and is the team leader.